arXiv:1603.07965v1 [cs.CV] 25 Mar 2016

Unsupervised Category Discovery via Looped Deep Pseudo-Task Optimization Using a Large Scale Radiology Image Database

Xiaosong Wang Le Lu Hoo-chang Shin Lauren Kim Isabella Nogues Jianhua Yao
Ronald Summers
Imaging Biomarkers and Computer-aided Detection Laboratory
Department of Radiology and Imaging Sciences
National Institutes of Health Clinical Center
10 Center Drive, Bethesda, MD 20892
xiaosong.wang,le.lu,hoo-chang.shin,lauren.kim,isabella.nogues,rms@nih.gov, jyao@cc.nih.gov

Abstract

Obtaining semantic labels on a large scale radiology image database (215,786 key images from 61,845 unique patients) is a prerequisite yet bottleneck to train highly effective deep convolutional neural network (CNN) models for image recognition. Nevertheless, conventional methods for collecting image labels (e.g., Google search followed by crowd-sourcing) are not applicable due to the formidable difficulties of medical annotation tasks for those who are not clinically trained. This type of image labeling task remains non-trivial even for radiologists due to uncertainty and possible drastic inter-observer variation or inconsistency.

In this paper, we present a looped deep pseudo-task optimization procedure for automatic category discovery of visually coherent and clinically semantic (concept) clusters. Our system can be initialized by domain-specific (CNN trained on radiology images and text report derived labels) or generic (ImageNet based) CNN models. Afterwards, a sequence of pseudo-tasks are exploited by the looped deep image feature clustering (to refine image labels) and deep CNN training/classification using new labels (to obtain more task representative deep features). Our method is conceptually simple and based on the hypothesized "convergence" of better labels leading to better trained CNN models which in turn feed more effective deep image features to facilitate more meaningful clustering/labels. We have empirically validated the convergence and demonstrated promising quantitative and qualitative results. Category labels of significantly higher quality than those in previous work are discovered. This allows for further investigation of the hierarchical semantic nature of the given large-scale radiology image database.

1. Introduction

The rapid and tremendous success of applying deep convolutional neural networks (CNNs) [27, 47, 52] to many challenging computer vision tasks derives from the accessibility of the well-annotated ImageNet [13,42] and PASCAL VOC [16] datasets. Deep CNNs perform significantly better than previous shallow learning methods and hand-crafted image features, however, at the cost of requiring greater amounts of training data. ImageNet pre-trained deep CNN models [22, 27, 32] serve an indispensable role to be bootstrapped upon for all externally-sourced data exploitation tasks [5, 31]. In the medical domain, however, no comparable labeled large-scale image dataset is available except the recent [44]. Vast amounts of radiology images/reports are stored in many hospitals' Picture Archiving and Communication Systems (PACS), but the main challenge lies in how to obtain ImageNet-level semantic labels on a large collection of medical images [44].

Nevertheless, conventional means of collecting image labels (e.g. Google image search using the terms from WordNet ontology hierarchy [34], SUN/PLACE databases [60, 63] or NEIL knowledge base [7]; followed by crowdsourcing [13]) are not applicable due to 1) the formidable difficulties of medical annotation tasks for clinically untrained annotators, 2) unavailability of a high quality or large capacity medical image search engine. On the other hand, even for well-trained radiologists, this type of "assigning labels to images" task is not aligned with their regular diagnostic routine work so that drastic inter-observer variations or inconsistency may be demonstrated. The protocols of defining image labels based on visible anatomic structures (often multiple), or pathological findings (possibly multiple) or using both cues have a lot of ambiguity.

Shin et al. [44] first extract the sentences depicting disease reference key images (similar concept to "key frames

in videos") using natural language processing (NLP) out of $\sim$ 780K patients' radiology reports, and find 215,786 key images of 61,845 unique patients from PACS. Then, image categorization labels are mined via unsupervised hierarchical Bayesian document clustering, i.e. generative latent Dirichlet allocation (LDA) topic modeling [3], to form 80 classes at the first level of hierarchy. The purely text-computed category information offers some coarse level of radiology semantics but is limited in two aspects: 1) The classes are *highly unbalanced*, in which one dominating category contains 113,037 images while other classes contain a few dozens. 2) The classes are not *visually* coherent. As a result, transfer learning from the CNN models trained in [44] to other medical computer-aided detection (CAD) problems performs less compellingly than those transferred directly from ImageNet CNNs [27, 46, 52].

In this paper, we present a **L**ooped **D**eep **P**seudo-task **O**ptimization (LDPO) approach for automatic category discovery of visually coherent and clinically semantic (concept) clusters. The true semantic category information is assumed to be latent and not directly observable. The main idea is to learn and train CNN models using pseudo-task labels (when human annotated labels are unavailable) and iterate this process with the expectation that pseudo-task labels will eventually resemble latent true image categories. Our work is partly related to the recent progress of semi-supervised learning or self-taught image classification, which has advanced both image classification and clustering processes [11, 12, 24, 30, 38, 48]. The iterative optimization in [24, 48] seeks to identify discriminative local visual patterns and reject others, whereas our goal is to set better labels for all images during iterations towards auto-annotation.

Our contributions are in several fold. **1)**, We propose a new "iteratively updated" deep CNN representation based on the LDPO technique. Thus it requires no hand-crafted image feature engineering [24, 30, 38, 48] which may be challenging for a large scale medical image database. Our method is conceptually simple and based on the hypothesized "convergence" of better labels lead to better trained CNN models which in turn, offer more effective deep image features to facilitate more meaningful clustering/labels. *This looped property is unique to deep CNN classification-clustering models since other types of classifiers do not learn better image features simultaneously.* We use the database from [44] to conduct experiments with the proposed method in different LDPO settings. Specifically, different pseudo-task initialization strategies, two CNN architectures of varying depths (i.e., AlexNet [27] and GoogLeNet [52]), different deep feature encoding schemes [8, 9] and clustering via K-means only or over-fragmented K-means followed by Regularized Information Maximization (RIM [20] as an effective model selection method), are extensively explored and empirically evaluated. **2)**, We consider the deep feature clustering followed by supervised CNN training as the outer loop and the deep feature clustering as the inner loop. Model selection on the number of clusters is critical and we carefully employ over-fragmented K-means followed by RIM model pruning/tuning to implement this criterion. This helps prevent cluster labeling amongst similar images, which can consequently compromise the CNN model training in the outer loop iteration. **3)**, The convergence of our LDPO framework can be observed and measured in both the cluster-similarity score plots and the CNN training classification accuracies. **4)**, Given the deep CNN LDPO models, hierarchical category relationships in a tree-like structure can be naturally formulated and computed from the final pairwise CNN classification confusion measures, as described in 3.5. We will make our discovered image annotations (after reviewed and verified by board-certified radiologists in a with-humans-in-the-loop fashion [62]) together with trained CNN models publicly available upon publication.

To the best of our knowledge, this is the first work exploiting to integrate unsupervised deep feature clustering and supervised deep label classification for self-annotating a large scale radiology image database where the conventional means of image annotation are not feasible. The measurable LDPO "convergence" makes this ill-posed problem well constrained, at no human labeling costs. Our proposed LDPO method is also quantitatively validated using Texture-25 dataset [12, 29] where the "unsupervised" classification accuracy improves over LDPO iterations. The ground truth labels of texture images [12, 29] are known and used to measure the accuracy scores against LDPO clustering labels. Our results may grant the possibility of 1), investigating the hierarchical semantic nature (object/organ, pathology, scene, modality, etc.) of categories [23, 40]; 2), finer level image mining for tag-constrained object instance discovery and detection [1, 59], given the large-scale radiology image database.

2. Related Work

Unsupervised and Semi-supervised Learning: Dai *et al.* [11, 12] study the semi-supervised image classification/clustering problem on texture [29], small to middle-scale object classes (e.g., Caltech-101 [17]) and scene recognition datasets [37]. By exploiting the data distribution patterns that are encoded by so called ensemble projection (EP) on a rich set of visual prototypes, the new image representation derived from clustering is learned for recognition. Graph based approaches [26, 33] are used to link the unlabeled image instances to labeled ones as anchors and propagate labels by exploiting the graph topology and connectiveness weights. In an unsupervised manner, Coates *et al.* [10] employ k-means to mine image patch filters and

then utilize the resulted filters for feature computation. Surrogate classes are obtained by augmenting each image patch with its geometrically transformed versions and a CNN is trained on top of these surrogate classes to generate features [15]. Wang *et al.* [56] design a Siamese-triplet CNN network, leveraging object tracking information in 100K unlabeled videos to provide the supervision for visual representation learning. Our work initializes an unlabeled image collection with labels from a pseudo-task (e.g., text topic modeling generated labels [44]) and update the labels through an iterative looped optimization of deep CNN feature clustering and CNN model training (towards better deep image features).

Text and Image: [28] is a seminal work that models the semantic connections between image contents and the text sentences. Those texts describe cues of detecting objects of interest, attributes and prepositions and can be applied as contextual regularizations. [25] proposes a structured objective to align the CNN based image region descriptors and bidirectional Recurrent Neural Networks (RNN) over sentences through the multimodal embedding. [55] presents a deep recurrent architecture from "Sequence to Sequence" machine translation [51] to generate image description in natural sentences, via maximizing the likelihood of the target description sentence given the training image. [49] applies extensive NLP parsing techniques (e.g., unigram terms and grammatical relations) to extract concepts that are consequently filtered by the discriminative power of visual cues and grouped by joint visual and semantic similarities. [6] further investigates an image/text co-clustering framework to disambiguate the multiple semantic senses of some Polysemy words. The NLP parsing in radiology reports is arguably much harder than processing those public datasets of image captions [25, 28, 55] where most plain text descriptions are provided. Radiologists often rule out or indicate pathology/disease terms, *not existing in the corresponding key images*, but based on patient priors and other long-range contexts or abstractions. In [45], only $\sim 8\%$ key images (18K out of 216K) can be tagged from NLP with the moderate confidence levels. We exploit the interactions from the text-derived image labels, to the proposed LDPO (mainly operating in the image modality) and the final term extraction from image groups.

Domain Transfer and Auto-annotation: Deep CNN representation has made transfer learning or domain adaption among different image datasets practical, via straightforward fine-tuning [19, 39]. Using pre-trained deep CNNs allows for the cross-domain transfer between weakly supervised video labels and noisy image labels. It can further output localized action frames by mutually filtering out low CNN-confidence instances [50]. A novel CNN architecture is exploited for deep domain transfer to handle unlabeled and sparsely labeled target domain data [54]. An image label auto-annotation approach is addressed via multiple instance learning [58] but the target domain is restricted to a small subset (25 out of 1000 classes) of ImageNet [13] and SUN [60]. [57] introduces a method to identify a hierarchical set of unlabeled data clusters (spanning a spectrum of visual concept granularities) that are efficiently labeled to produce high performing classifiers (thus less label noise at instance level). By learning visually coherent and class balanced labels through LDPO, we expect that the studied large-scale radiology image database can markedly improve its feasibility in domain transfer to specific CAD problems where very limited training data are available per task.

3. Looped Deep Pseudo-Task Optimization

Traditional detection and classification problems in medical imaging, e.g. Computer Aided Detection (CAD) [41], require precise labels of lesions or diseases as the training/testing ground-truth. This usually requires a large amount of annotation from well-trained medical professionals (especially at the era of "deep learning"). Employing and converting the medical records stored in the PACS into labels or tags is very challenging [45]. Our approach performs the category discovery in an empirical manner and returns accurate key-word category labels for all images, through an iterative framework of deep feature extraction, clustering, and deep CNN model fine-tuning.

As illustrated in Fig. 1, the iterative process begins by extracting the deep CNN feature based on either a fine-tuned (with high-uncertainty radiological topic labels [44]) or generic (from ImageNet labels [27]) CNN model. Next, the deep feature clustering with k-means or k-means followed by RIM is exploited. By evaluating the purity and mutual information between discovered clusters, the system either terminates the current iteration (which leads to an optimized clustering output) or takes the refined cluster labels as the input to fine-tune the CNN model for the following iteration. Once the visually coherent image clusters are obtained, the system further extracts semantically meaningful text words for each cluster. All corresponding patient reports per category cluster are finally adopted for the NLP. Furthermore, the hierarchical category relationship is built using the class confusion measures of the latest converged CNN classification models.

3.1. Convolution Neural Networks

The proposed LDPO framework is applicable to a variety of CNN models. We analyze the CNN activations from layers of different depths in AlexNet [27] and GoogLeNet [52]. Pre-trained models on the ImageNet ILSVRC data are obtained from Caffe Model Zoo [22]. We also employ the Caffe CNN implementation [22] to perform fine-tuning on pre-trained CNNs using the key image database (from [44]). Both CNN models with/without fine-tuning are used to

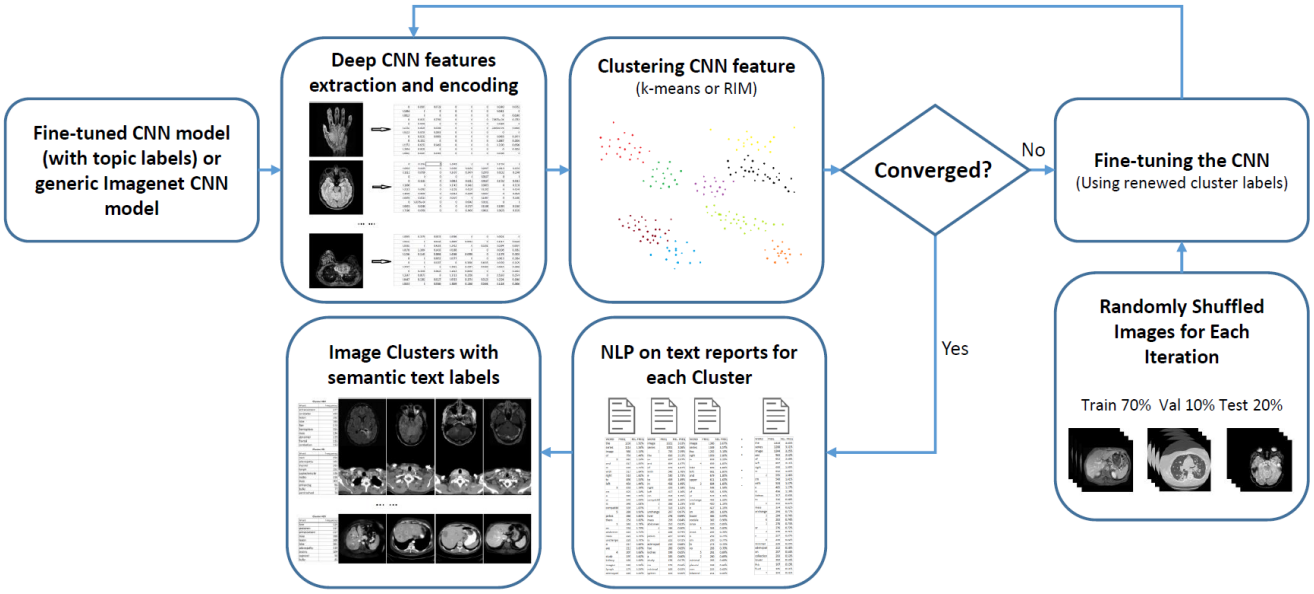


Figure 1. The overview of looped deep pseudo-task optimization framework.

Table 1. Configurations of CNN output layers and encoding methods (Output dimension is 4096, except the last row as 1024).

CNN model	Layer	Activations	Encoding
AlexNet	Conv5	$(13, 13, 256)$	FV+PCA
AlexNet	Conv5	$(13, 13, 256)$	VLAD+PCA
AlexNet	FC7	4096	–
GoogLeNet	Inc.5b	$(7, 7, 1024)$	VLAD+PCA
GoogLeNet	Pool5	1024	–

initialize the looped optimization. AlexNet is a common CNN architecture with 7 layers and the extracted features from its convolutional or fully-connected layers have been broadly investigated [19,25,39]. The encoded convolutional features for image retrieval tasks are introduced in [35], which verifies the image representation power of convolutional features. In our experiments we adopt feature activations of both the 5th convolutional layer $Conv5$ and 7th fully-connected (FC) layer $FC7$ as suggested in [4, 9]. GoogLeNet is a much deeper CNN architecture compared to AlexNet, which comprises 9 inception modules and an average pooling layer. Each inception modules is truly a set of convolutional layers with multiple window sizes, i.e. $1 \times 1, 3 \times 3, 5 \times 5$. Similarly, we explore the deep image features from the last inception layer $Inception5b$ and final pooling layer $Pool5$. Table 1 illustrates the detailed model layers and their activation dimensions.

3.2. Encoding Images using Deep CNN Features

While the features extracted from fully-connected layer are able to capture the overall layout of objects inside the image, features computed at the last convolution layer preserve the local activations of images. Different from the standard max-pooling before feeding the fully-connected layer, we adopt the same setting ([8]) to encode the convolutional layer outputs in a form of dense pooling via Fisher Vector (FV) [36] and Vector Locally Aggregated Descriptor (VLAD) [21]. Nevertheless, the dimensions of encoded features are much higher than those of the FC feature. Since there is redundant information from the encoded features and we intend to make the results comparable between different encoding schemes, Principal Component Analysis (PCA) is performed to reduce the dimensionality to 4096, equivalent to the FC features' dimension.

Given a pre-trained (generic or domain-specific) CNN model (i.e., Alexnet or GoogLeNet), an input image I is resized to fit the model definition and feed into the CNN model to extract features $\{f^L_{i,j}\}$ $(1 \leqslant i, j \leqslant s^L)$ from the L-th convolutional layer with dimensions $s^L \times s^L \times d^L$, e.g., $13 \times 13 \times 256$ of $Conv5$ in AlexNet and $7 \times 7 \times 1024$ of $Pool5$ in GoogLeNet. For Fisher Vector implementation, we use the settings as suggested in [9]: 64 Gaussian components are adopted to train the Gaussian mixture Model(GMM). The dimension of resulted FV features is significantly higher than $FC7$'s, i.e. $32768(2 \times 64 \times 256)$ vs 4096. After PCA, the FV representation per image is reduced to a 4096-component vector. A list of deep image features, the encoding methods and output dimensions are provided in Table 1. To be consistent with the settings of FV representation, we initialize the VLAD encoding of convolutional image features by k-means clustering with $k = 64$. Thus the dimensions of VLAD descriptors are $16384(64 \times 256)$ of $Conv5$ in AlexNet and $65536(64 \times 1024)$ of $Inception5b$ in GoogLeNet. PCA further reduces the dimensions of both to 4096.

3.3. Image Clustering

Image clustering plays an indispensable role in our LDPO framework. We hypothesize that the newly generated clusters driven by looped pseudo-task optimization are better than the previous ones in the following terms: 1) Images in each cluster are visually more coherent and discriminative from instances in other clusters; 2) The numbers of images per cluster are approximately equivalent to achieve class balance; 3) The number of clusters is self-adaptive according to the statistical properties of a large collection of image data. Two clustering methods are employed here, i.e. k-means alone and an over-segmented k-means (where K is much larger than the first setting, e.g., 1000) followed by Regularized Information Maximization (RIM) [20] for model selection and optimization.

k-means is an efficient clustering algorithm provided that the number of clusters is known. We explore k-means clustering here for two reasons: 1) To set up the baseline performance of clustering on deep CNN image features by fixing the number of clusters k at each iteration; 2) To initialize the RIM clustering since k-means is only capable of fulfilling our first two hypotheses, and RIM will help satisfy the third. Unlike k-means, RIM works with fewer assumptions on the data and categories, e.g. the number of clusters. It is designed for discriminative clustering by maximizing the mutual information between data and the resulted categories via a complexity regularization term. The objective function is defined as

$$f(\mathbf{W};\mathbf{F},\lambda) = I_{\mathbf{W}}\{c;\mathbf{f}\} - R(\mathbf{W};\lambda), \tag{1}$$

where $c \in \{1,...,K\}$ is a category label, $\mathbf{F}$ is the set of image features $\mathbf{f_i} = (f_{i1},...,f_{iD})^T \in \mathbb{R}^D$. $I_{\mathbf{W}}\{c;\mathbf{f}\}$ is an estimation of the mutual information between the feature vector $\mathbf{f}$ and the label c under the conditional model $p(c|\mathbf{f},\mathbf{W})$. $R(\mathbf{W};\lambda)$ is the complexity penalty and specified according to $p(c|\mathbf{f},\mathbf{W})$. As demonstrated in [20], we adopt the unsupervised multilogit regression cost. The conditional model and the regularization term are consequently defined as

$$p(c=k|\mathbf{f},\mathbf{W}) \propto exp(w_k^T\mathbf{f}+b_k) \tag{2}$$

$$R(\mathbf{W};\lambda) = \lambda\sum_k w_k^T w_k, \tag{3}$$

where $\mathbf{W} = \{\mathbf{w}_1,...,\mathbf{w}_K,b_1,...,b_K\}$ is the set of parameters and $\mathbf{w}_k \in \mathbb{R}^D, b_k \in \mathbb{R}$. Maximizing the objective function is now equivalent to solving a logistic regression problem. R is the L_2 regulator of weight $\{w_k\}$ and its power is controlled by λ. Large λ values will enforce to reduce the total number of categories considering that no penalty is given for unpopulated categories [20]. This characteristic enables RIM to attain the optimal number of categories coherent to the data. λ is fixed to 1 in all our experiment.

3.4. Convergence in Clustering and Classification

Before exporting the newly generated cluster labels to fine-tune the CNN model of the next iteration, the LDPO framework will evaluate the quality of clustering to decide if convergence has been achieved. Two convergence measurements have been adopted [53], i.e., Purity and Normalized Mutual Information (NMI). We take these two criteria as forms of empirical similarity examination between two clustering results from adjacent iterations. If the similarity is above a certain threshold, we believe the optimal clustering-based categorization of the data is reached. We indeed find that the final number of categories from the RIM process in later LDPO iterations stabilize around a constant number. The convergence on classification is directly observable through the increasing top-1, top-5 classification accuracy levels in the initial few LDPO rounds which eventually fluctuate slightly at higher values.

Convergence in clustering is achieved by adopting the underlying classification capability stored in those deep CNN features through the looped optimization, which accents the visual coherence amongst images inside each cluster. Nevertheless, the category discovery of medical images will further entail clinically semantic labeling of the images. From the optimized clusters, we collect the associated text reports for each image and assemble each cluster's text reports together as a unit. Then NLP is performed on each report unit to find highly recurring words to serve as key word labels for each cluster by simply counting and ranking the frequency of each word. Common words to all clusters are removed from the list. The resultant key words and randomly sampled exemplary images are ultimately compiled for review by board-certified radiologists. This process shares some analogy to the human-machine collaborated image database construction [57, 62]. In future work, NLP parsing (especially term negation/assertion) and clustering can be integrated into LDPO framework.

3.5. Hierarchical Category Relationship

ImageNet [13] are constructed according to WordNet ontology hierarchy [34]. Recently, a new formalism so-called Hierarchy and Exclusion (HEX) graphs has been introduced [14] to perform object classification by exploiting the rich structure of real world labels [13, 27]. In this work, our converged CNN classification model can be further extended to explore the hierarchical class relationship in a tree representation. First, the pairwise class similarity or affinity score $A_{i,j}$ between class (i,j) is modeled via an adapted measurement from CNN classification confusion [5].

$$A_{i,j} = \frac{1}{2}\Big(Prob(i|j) + Prob(j|i)\Big) \tag{4}$$

$$= \frac{1}{2}\Big(\frac{\sum_{I_m\in C_i} CNN(I_m|j)}{|C_i|} + \frac{\sum_{I_n\in C_j} CNN(I_n|i)}{|C_j|}\Big) \tag{5}$$

where C_i, C_j are the image sets for class i,j respectively, $|\cdot|$ is the cardinality function, $CNN(I_m|j)$ is the CNN classification score of image I_m from class C_i at class j obtained directly by the N-way CNN flat-softmax. Here $A_{i,j} = A_{j,i}$ is symmetric by averaging $Prob(i|j)$ and $Prob(j|i)$.

Affinity Propagation algorithm [18] (AP) is invoked to perform "tuning parameter-free" clustering on this pairwise affinity matrix $\{A_{i,j}\} \in \mathbb{R}^{K\times K}$. This process can be executed recursively to generate a hierarchically merged category tree. Without loss of generality, we assume that at level L, classes i^L,j^L are formed by merging classes at level L-1 through AP clustering. The new

affinity score A_{i^L,j^L} is computed as follows.

$$A_{i^L,j^L} = \frac{1}{2}\Big(Prob(i^L|j^L) + Prob(j^L|i^L)\Big) \quad (6)$$

$$Prob(i^L|j^L) = \frac{\sum_{I_m \in C_{i^L}} \sum_{k \in j^L} CNN(I_m|k)}{|C_{i^L}|} \quad (7)$$

where L-th level class label j^L include all merged original classes (i.e., 0-th level before AP is called) $k \in j^L$ so far. From the above, the N-way CNN classification scores (Sec. 3.4) only need to be evaluated once. A_{i^L,j^L} at any level can be computed by summing over these original scores. The discovered category hierarchy can help alleviate the highly uneven visual separability between different object categories in image classification [61] from which the category-embedded hierarchical deep CNN could be beneficial.

4. Experimental Results & Discussion

4.0.1 Dataset:

We experiment on the same dataset used in [44]. The image database contains totally 216K $2D$ key-images which are associated with $\sim$ 62K unique patients' radiology reports. Key-images are directly extracted from the Dicom file and resized as 256×256 bitmap images. Their intensity ranges are rescaled using the default window settings stored in the Dicom header files (this intensity rescaling factor improves the CNN classification accuracies by $\sim 2\%$ to [44]). Linked radiology reports are also collected as separate text files with patient-sensitive information removed for privacy reasons. At each LDPO iteration, the image clustering is first applied on the entire image dataset so that each image will receive a cluster label. Then the whole dataset is randomly reshuffled into three subgroups for CNN fine-tuning via Stochastic Gradient Descent (SGD): i.e. training (70%), validation (10%) and testing (20%). In this way, the convergence is not only achieved on a particular data-split configuration but generalized to the entire database.

In order to quantitatively validate our proposed LDPO framework, we also apply category discovery on the texture-25 dataset [12, 29]: 25 texture classes, with 40 samples per class. The images from Texture-25 appear drastically different from those natural images in ImageNet, similar to our domain adaptation task from natural to radiology images. The ground truth labels are first hidden from the unsupervised LDPO learning procedure and then revealed to produce the quantitative measures (where purity becomes accuracy) against the resulted clusters. The cluster number is assumed to be known to LDPO and thus the model selection module of RIM in clustering is dropped.

4.0.2 CNN Fine-tuning:

The Caffe [22] implementation of CNN models are used in the experiment. During the looped optimization process, the CNN is fine-tuned for each iteration once a new set of image labels is generated from the clustering stage. Only the last softmax classification layer of the models (i.e. 'FC8' in AlexNet and 'loss3/classifier' in GoogLeNet) is significantly modulated by 1) setting a higher learning rate than all other layers and 2) updating the (varying but converging) number of category classes from the newly computed results of clustering.

4.1. LDPO Convergence Analysis

We first study how the different settings of proposed LDPO framework will affect convergence as follows:

4.1.1 Clustering Method:

We perform k-means based image clustering with $k \in \{80, 100, 200, 300, 500, 800\}$. Fig. 2 shows the changes of top-1 accuracy, cluster purity and NMI with different k across iterations. The classification accuracies quickly plateau after 2 or 3 iterations. Smaller k values naturally trigger higher accuracies ($> 86.0\%$ for $k = 80$) as less categories make the classification task easier. Levels of Purity and NMI between clusters from two consecutive iterations increase quickly and fluctuate close to 0.7, thus indicating the convergence of clustering labels (and CNN models). The minor fluctuation are rather due to the randomly re-sorting of the dataset in each iteration. RIM clustering takes an over-segmented k-means results as initialization, e.g., $k = 1000$ in our experiments. As shown in Fig. 3 Top-left, RIM can estimate the category capacities or numbers consistently under different image representations (deep CNN feature + encoding approaches). k-means clustering enables LDPO to approach the convergence quickly with high classification accuracies; whereas, the added RIM based model selection delivers more balanced and semantically meaningful clustering results (see more in Sec. 4.2). This is due to RIM's two unique characteristics: 1), less restricted geometric assumptions in the clustering feature space; 2), the capacity to attain the optimal number of clusters by maximizing the mutual information of input data and the induced clusters via a regularized term.

4.1.2 Pseudo-Task Initialization:

Both ImageNet and domain-specific [44] CNN models have been employed to initialize the LDPO framework. In Fig. 3, two CNNs of AlexNet-FC7-ImageNet and AlexNet-FC7-Topic demonstrate their LDPO performances. LDPO initialized by ImageNet CNN reach the steady state noticeably slower than its counterpart, as AlexNet-FC7-Topic already contains the domain information from this radiology image database. However, similar clustering outputs are produced after convergence. Letting LDPO reach $\sim$ 10 iterations, two different initializations end up with very close clustering results (i.e., Cluster number, purity and NMI) and similar classification accuracies (shown in Table 2).

4.1.3 CNN Deep Feature and Image Encoding:

Different image representations can vary the performance of proposed LDPO framework as shown in Fig. 3. As mentioned in Sec. 3, deep CNN images features extracted from different layers of CNN models (AlexNet and GoogLeNet) contain the level-specific visual information. Convolutional layer features retain the spatial activation layouts of images while FC layer features do not. Different encoding approaches further lead to various outcomes of our LDPO framework. The numbers of clusters range from 270 (**AlexNet-FC7-Topic** with no deep feature encoding) to 931 (the more sophisticated **GoogLeNet-Inc.5b-VLAD** with VLAD encoding). The numbers of clusters discovered by RIM re-

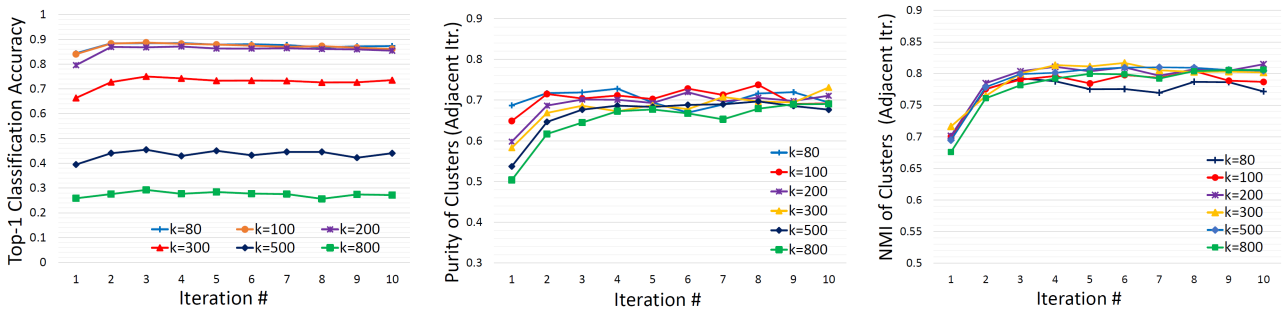


Figure 2. Performance of LDPO using k-means clustering with a variety of k. From left to right, the top-1 classification accuracy and the purity and NMI of clusters from adjacent iterations are shown.

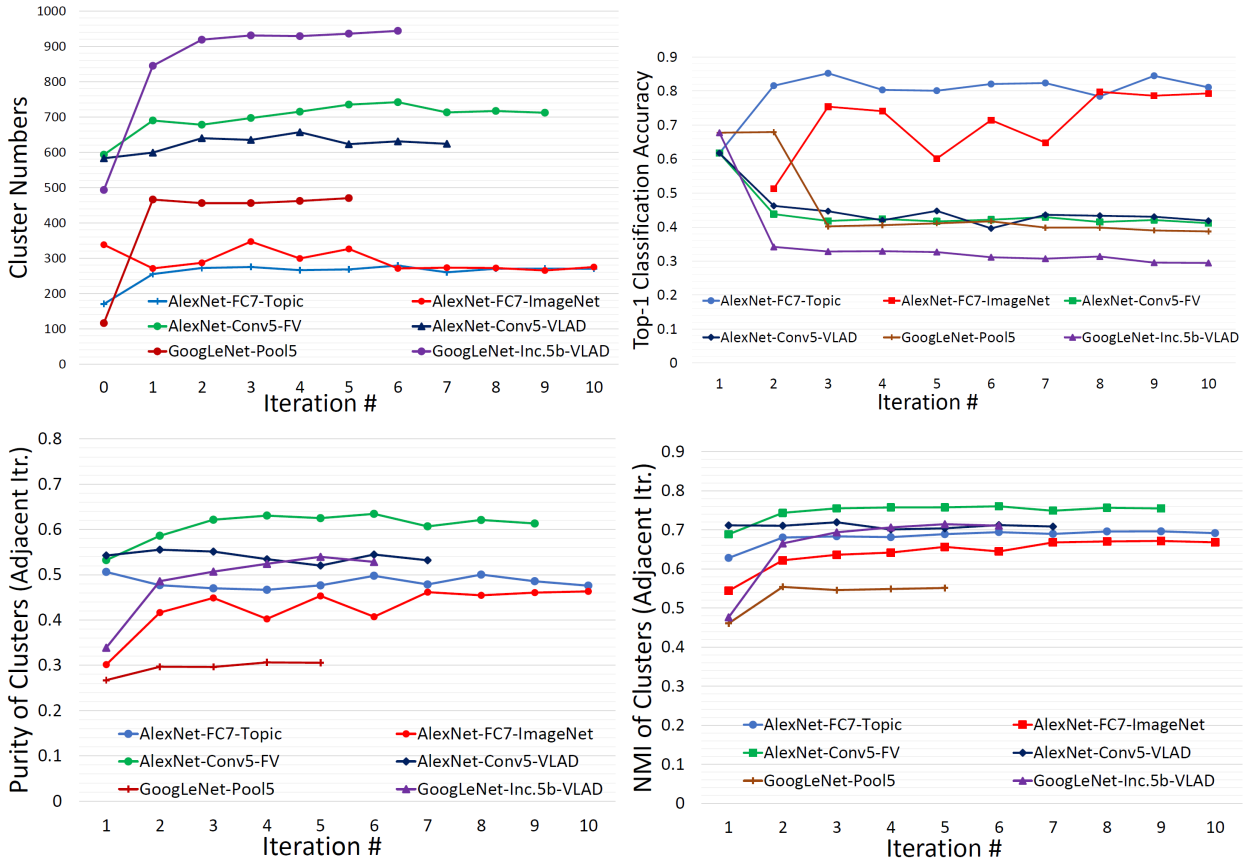


Figure 3. Performance of LDPO using RIM clustering with different image encoding methods (*i.e.*, FV and VLAD) and CNN Architectures (*i.e.*, AlexNet and GoogLeNet). From left to right(top to bottom), the number of clusters discovered, Top-1 accuracy of trained CNNs, the purity and NMI of clusters from adjacent iterations are illustrated.

flect the amount of information complexity stored in the radiology database.

4.1.4 Computational Cost:

LDPO runs on a node of Linux computer cluster with 16 CPU cores (x2650), 128G memory and Nvidia K20 GPUs. The Computational costs of different LDPO configurations are shown in Table 2 per looped iteration. The more sophisticated and feature rich settings, e.g., **AlexNet-Conv5-FV**, **GoogLeNet-Pool5** and **GoogLeNet-Inc.5b-VLAD**, require more time to converge.

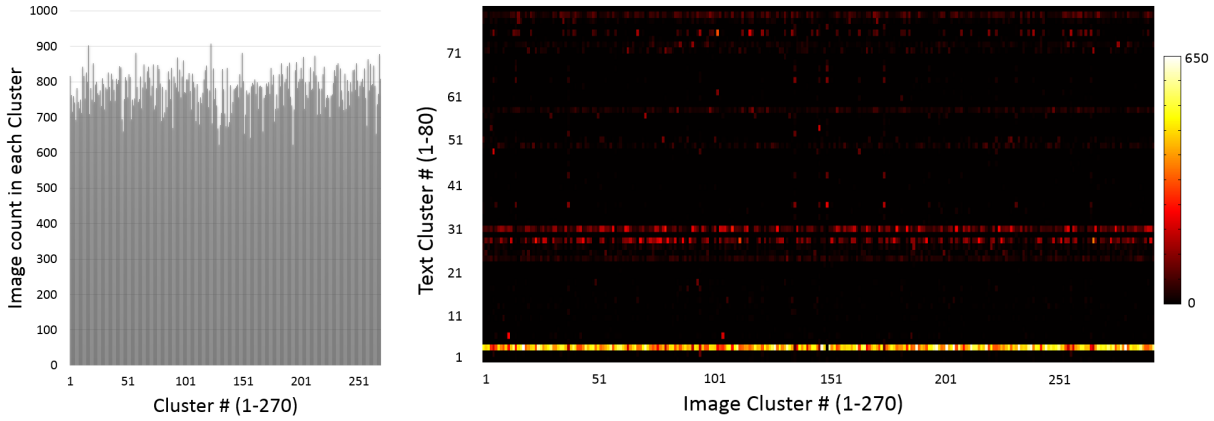


Figure 4. Statistics of converged categories using the Alexnet-FC7-Topic setting. Left: the image numbers in each cluster; Right: affinity matrix of two clustering results (AlexNet-FC7-270 vs Text-Topics-80 produced using the approach in [44]).

Table 2. Classification Accuracy of Converged CNN Models

CNN setting	Cluster #	Top-1	Top-5
AlexNet-FC7-Topic	270	0.8109	0.9412
AlexNet-FC7-ImageNet	275	0.8099	0.9547
AlexNet-Conv5-FV	712	0.4115	0.4789
AlexNet-Conv5-VLAD	624	0.4333	0.5232
GoogLeNet-Pool5	462	0.4109	0.5609
GoogLeNet-Inc.5b-VLAD	929	0.3265	0.4001

Table 3. Computational Cost of LDPO

CNN setting	Time per iter.(HH:MM)
AlexNet-FC7-Topic	14:35
AlexNet-FC7-Imagenet	14:40
AlexNet-Conv5-FV	17:40
AlexNet-Conv5-VLAD	15:44
GoogLeNet-Pool5	21:12
GoogLeNet-Inc.5b-VLAD	23:35

4.2. LDPO Categorization and Auto-annotation Results

The category discovery clusters employing our LDPO method are found to be more visually coherent and cluster-wise balanced in comparison to the results in [44] where clusters are formed only from text information ($\sim 780K$ radiology reports). Fig. 4 **Left** shows the image numbers for each cluster from the AlexNet-FC7-Topic setting. The numbers are uniformly distributed with a mean of 778 and standard deviation of 52. Fig. 4 **Right** illustrates the relation of clustering results derived from image cues or text reports [44]. Note that there is no instance-balance-per-cluster constraints in the LDPO clustering. The clusters in [44] are highly uneven: 3 clusters inhabit the majority of images. Fig. 5 shows sample images and top-10 associated key words from 4 randomly selected clusters (more results in the supplementary material). The LDPO clusters are found to be semantically or clinically related to the corresponding key words, containing the information of (likely appeared) anatomies, pathologies (e.g., adenopathy, mass), their attributes (e.g., bulky, frontal) and imaging protocols or properties.

Next, from the best performed LDPO models in Table 2, **AlexNet-FC7-Topic** has **Top-1** classification accuracy of 0.8109 and **Top-5** accuracy 0.9412 with 270 formed image categories; **AlexNet-FC7-ImageNet** achieves accuracies of 0.8099 and 0.9547, respectively, from 275 discovered classes. In contrast, [44] reports **Top-1** accuracies of 0.6072, 0.6582 and **Top-5** as 0.9294, 0.9460 on 80 text only computed classes using AlexNet [27] or VGGNet-19 [47], respectively. Markedly better accuracies (especially on **Top-1**) on classifying higher numbers of classes (being generally more difficult) highlight advantageous quality of the LDPO discovered image clusters or labels. This means that the LDPO results have rendered significantly better performance on automatic image labeling than the most related previous work [44], under the same radiology database. *After the subjective evaluation by two board-certified radiologists,* **AlexNet-FC7-Topic** *of 270 categories and* **AlexNet-FC7-ImageNet** *of 275 classes are preferred, out of total six model-encoding setups. Interestingly, both CNN models have no deep feature encoding built-in and preserve the gloss image layouts (capturing somewhat global visual scenes without unordered FV or VLAD encoding schemes [8, 9, 21].).*

For the quantitative validation, LDPO is also evaluated on the Texture-25 dataset as an unsupervised texture classification problem. The purity and NMI are computed between the resulted LDPO clusters per iteration and the ground truth clusters (of 25 texture image classes [12,29]) where purity becomes classification accuracy. **AlexNet-FC7-ImageNet** is employed and the quantitative results are plotted in Fig. 7. Using the same clustering method of k-means, the purity or accuracy measurements improve from 53.9% (0-th) to 66.1% at the 6-th iteration, indicating that LDPO indeed learns better deep image features and labels in the looped process. Similar trend is found for another texture dataset [8]. Exploiting LDPO for other domain transfer based auto-annotation

Cluster #5

Word	Frequency
neck	656
adenopathy	343
thyroid	295
lymph	292
supraclavicular	236
nodes	218
mass	203
enhancing	96
bulky	77
paratracheal	76

Cluster #14

Word	Frequency
calf	369
mass	263
subcutaneous	205
thigh	204
lesion	127
lower	124
enhancing	111
bone	105
fossa	92
nerve	88

Cluster #23

Word	Frequency
liver	524
abdomen	337
enhancement	217
mass	198
lesion	168
lobe	161
adenopathy	119
lesions	109
segment	58
bulky	45

Cluster #64

Word	Frequency
enhancement	277
cerebellar	193
lesion	192
lobe	186
flair	173
hemisphere	155
mass	134
abnormal	119
frontal	115
cerebellum	113

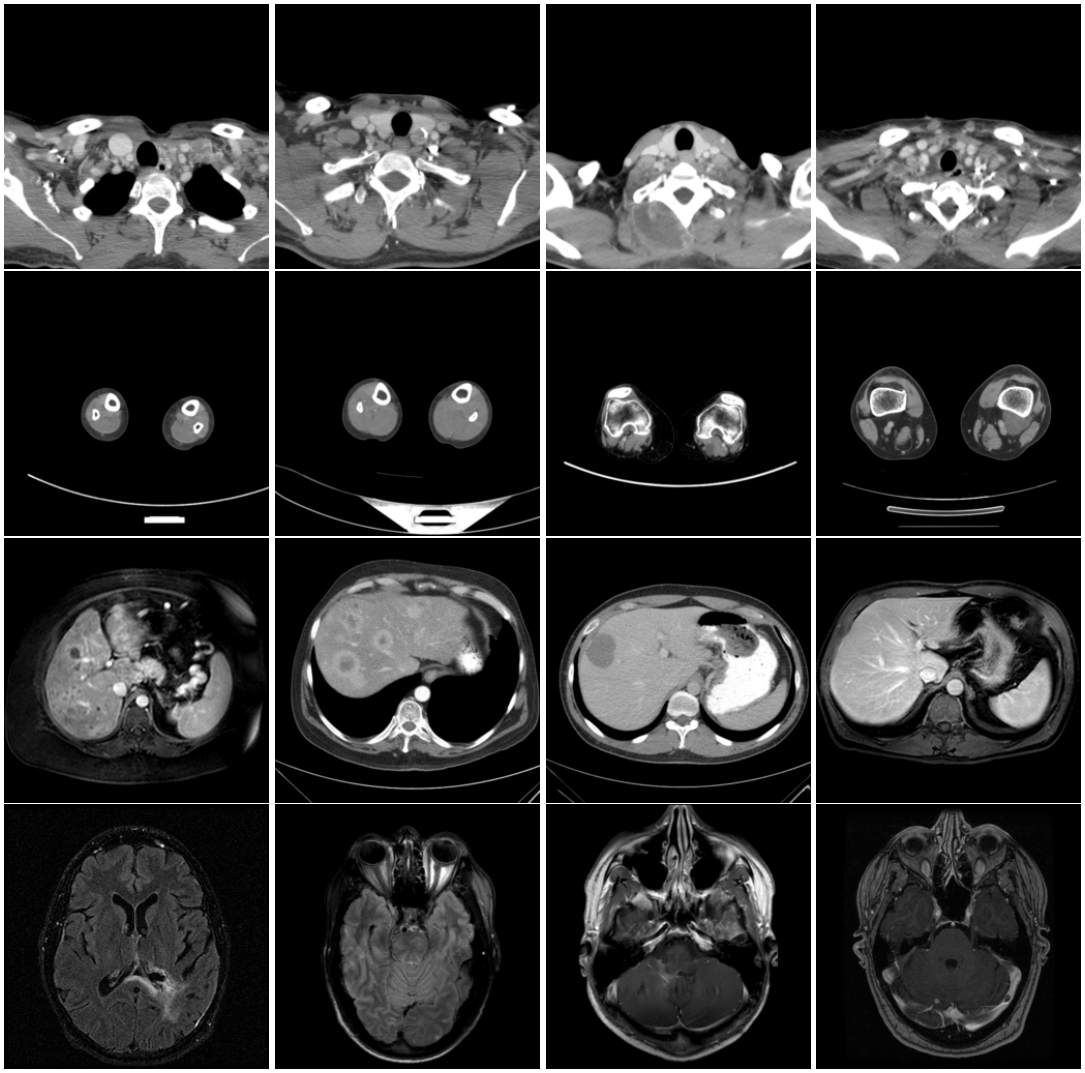


Figure 5. Sample images of four LDPO clusters with associated clinically semantic key words, containing the information of (likely appeared) anatomies, pathologies, their attributes and imaging protocols or properties.

tasks will be left as future work.

The final trained CNN classification models allow to compute the pairwise category similarities or affinity scores using the CNN classification confusion values between any pair of classes (Sec. 3.5). Affinity Propagation algorithm is called recursively to form a hierarchical category tree. The resulted category tree has (270, 64, 15, 4, 1) different class labels from bottom (leaf) to top (root). The random color coded category tree is shown in Fig. 6. The high majority of images in the clusters of this branch are verified as CT Chest scans by radiologists. *Enabling to construct a semantic and meaningful hierarchy of classes offers another indicator to validate the proposed LDPO category discovery method and results.* Refer to the supplementary material for more results. We will make our trained CNN models, computed deep image features and labels publicly available upon publication.

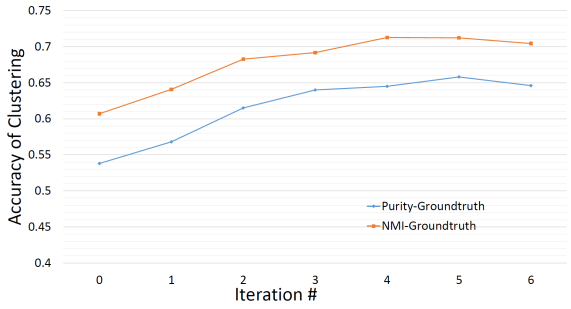


Figure 7. Purity (Accuracy) and NMI plots between the ground truth classes and LDPO discovered clusters versus the iteration numbers.

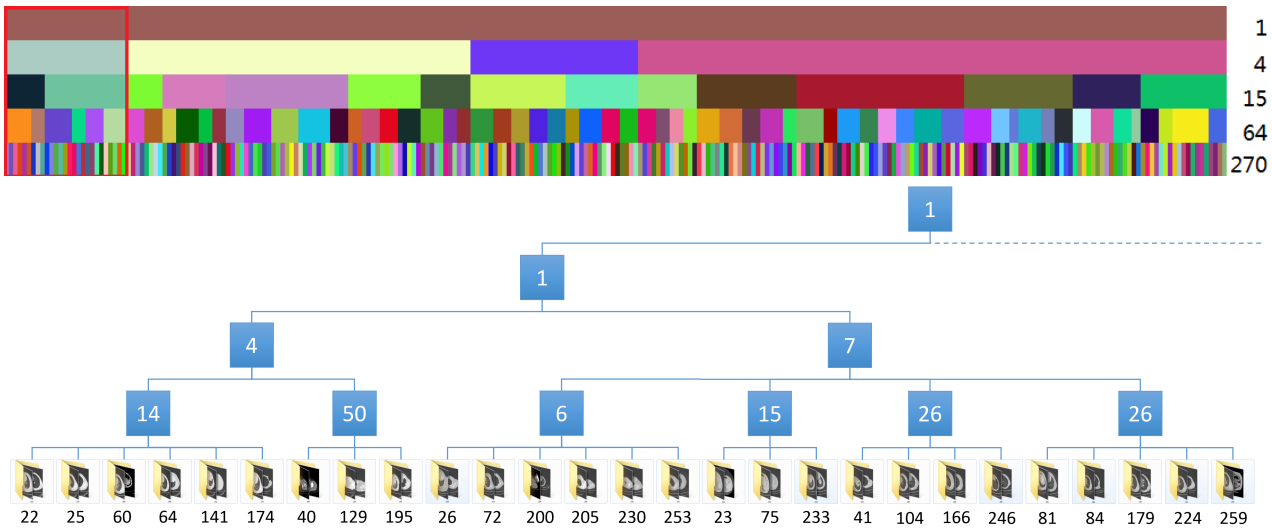


Figure 6. Five-level hierarchical categorization is illustrated with a randomized color for each cluster. Sample images and detailed tree structures from a branch (highlighted with a red bounding box) are also shown. *The high majority of images in the clusters of this branch are verified as CT Chest scans by radiologists*.

5. Conclusion & Future Work

In this paper, we present a new Looped Deep Pseudo-task Optimization framework to extract visually more coherent and semantically more meaningful categories from a large scale medical image database. We systematically and extensively conduct experiments under different settings of the LDPO framework to validate and evaluate its quantitative and qualitative performance. The measurable LDPO "convergence" makes the ill-posed auto-annotation problem well constrained without the burden of human labeling costs. For future work, we intend to explore the feasibility/performance on implementing our current LDPO clustering component by deep generative density models [2, 26, 43]. It may therefore be possible that both classification and clustering objectives can be built into a multi-task CNN learning architecture which is "end-to-end" trainable by alternating two task/cost layers during SGD optimization [54].

References

[1] L. Bazzani, A. Bergamo, D. Anguelov, and L. Torresani. Self-taught object localization with deep networks. *arXiv preprint arXiv:1409.3964*, 2015.

[2] Y. Bengio, I. Goodfellow, and A. Courville. Deep learning. Book in preparation for MIT Press, 2015.

[3] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3:993–1022, 2003.

[4] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *British Machine Vision Conference*, 2014.

[5] X. Chen and A. Gupta. Webly supervised learning of convolutional networks. In *Proc. of ICCV*, 2015.

[6] X. Chen, A. Ritter, A. Gupta, and T. Mitchell. Sense discovery via co-clustering on images and text. *Proc. of IEEE CVPR*, pages 5298–5306, 2015.

[7] X. Chen, A. Shrivastava, and A. Gupta. Neil: Extracting visual knowledge from web data. In *Proc. of ICCV*, 2013.

[8] M. Cimpoi, S. Maji, I. Kokkinos, and A. Vedaldi. Deep filter banks for texture recognition, description, and segmentation. *arXiv preprint arXiv:1507.02620*, 2015.

[9] M. Cimpoi, S. Maji, and A. Vedaldi. Deep filter banks for texture recognition and segmentation. *Proc. of IEEE CVPR*, 2015.

[10] A. Coates, A. Ng, and H. Lee. An analysis of single-layer networks in unsupervised feature learning. *AI and Statistics*, 2011.

[11] D. Dai and L. Van Gool. Ensemble projection for semi-supervised image classification. In *Proc. of ICCV*, 2013.

[12] D. Dai and L. Van Gool. Unsupervised high-level feature learning by ensemble projection for semi-supervised image classification and image clustering. Technical report, arXiv:1602.00955, 2016.

[13] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, pages 248–255. IEEE, 2009.

[14] K. Deng, N. Ding, Y. Jia, A. Frome, K. Murphy, S. Bengio, Y. Li, H. Neven, and H. Adam. Large-scale object classification using label relation graphs. *Proc. of ECCV*, pages 48–64, 2014.

[15] A. Dosovitskiy, J. Springenberg, M. Riedmiller, and T. Brox. Discriminative unsupervised feature learning with convolutional neural networks. *NIPS*, 2014.

[16] M. Everingham, S. M. A. Eslami, L. Van Gool, C. Williams, J. Winn, and A. Zisserman. The pascal visual object classes

challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015.

[17] L. Fei-Fei, R. Fergus, and P. Perona. Learning generative visual models from few training examples: an incremental bayesian approach tested on 101 object categories. *Proc. of IEEE CVPR workshop*, 2004.

[18] B. Frey and D. Dueck. Clustering by passing messages between data points. *Science*, 315:972–976, 2007.

[19] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Region-based convolutional networks for accurate object detection and semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2015.

[20] R. Gomes, A. Krause, and P. Perona. Discriminative clustering by regularized information maximization. *NIPS*, 2010.

[21] H. Jegou, F. Perronnin, M. Douze, J. Sanchez, P. Perez, and C. Schmid. Aggregating local image descriptors into compact codes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(9):1704–1716, Sept 2012.

[22] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.

[23] J. Johnson, R. Krishna, M. Stark, L. Li, D. Shamma, M. Bernstein, and L. Fei-Fei. Image retrieval using scene graphs. *Proc. of IEEE CVPR*, pages 3668–3678, 2015.

[24] M. Juneja, A. Vedaldi, C. Jawahar, and A. Zisserman. Blocks that shout: Distinctive parts for scene classification. *CVPR*, pages 923–930, 2013.

[25] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. *Proc. of IEEE CVPR*, pages 3128–3137, 2015.

[26] D. Kingma, S. Mohamed, D. Rezende, and M. Welling. Semi-supervised learning with deep generative models. *NIPS*, 2014.

[27] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[28] G. Kulkarni, V. Premraj, V. Ordonez, S. Dhar, S. Li, Y. Choi, A. Berg, and T. Berg. Babytalk: Understanding and generating simple image descriptions. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(12):2891–2903, 2013.

[29] S. Lazebnik, C. Schmid, and J. Ponce. A sparse texture representation using local affine regions. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(8):1265–1278, 2005.

[30] Y. Li and Z. Zhou. Towards making unlabeled data never hurt. *ICML*, 2011.

[31] X. Liang, S. Liu, Y. Wei, L. Liu, L. Lin, and S. Yan. Computational baby learning. In *Proc. of ICCV*, 2015.

[32] M. Lin, Q. Chen, and S. Yan. Network in network. In *Proc. of ICLR*, 2015.

[33] W. Liu, J. He, and S. Chang. Large graph construction for scalable semi-supervised learning. *ICML*, 2010.

[34] G. A. Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 1995.

[35] J. Y. Ng, F. Yang, and L. S. Davis. Exploiting local features from deep networks for image retrieval. *CoRR*, abs/1504.05133, 2015.

[36] F. Perronnin, J. Snchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In K. Daniilidis, P. Maragos, and N. Paragios, editors, *Computer Vision ECCV 2010*, volume 6314 of *Lecture Notes in Computer Science*, pages 143–156. Springer Berlin Heidelberg, 2010.

[37] A. Quattoni and A. Torralba. Recognizing indoor scenes. *Proc. of IEEE CVPR*, 2009.

[38] R. Raina, A. Battle, H. Lee, B. Packer, and A. Ng. Self-taught learning: transfer learning from unlabeled data. *ICML*, 2007.

[39] A. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. *ArXiv:1403.6382*, 2014.

[40] K. Rematas, B. Fernando, F. Dellaert, and T. Tuytelaars. Dataset fingerprints: Exploring image collections through data mining. *Proc. of IEEE CVPR*, pages 4867–4875, 2015.

[41] H. Roth, L. Lu, J. Liu, J. Yao, A. Seff, K. Cherry, E. Turkbey, and R. Summers. Improving computer-aided detection using convolutional neural networks and random view aggregation. In *IEEE Trans. on Medical Imaging*. 2015.

[42] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *arXiv preprint arXiv:1409.0575*, 2014.

[43] R. Salakhutdinov. Learning deep generative models. *Annual Review of Statistics and Its Application*, 2:361–385, 2015.

[44] H. Shin, L. Lu, L. Kim, A. Seff, J. Yao, and R. Summers. Interleaved text/image deep mining on a large-scale radiology database. *Proc. of IEEE CVPR*, 2015.

[45] H. Shin, L. Lu, L. Kim, A. Seff, J. Yao, and R. Summers. Interleaved text/image deep mining on a large-scale radiology image database for automated image interpretation. *arXiv:1505.00670*, 2015.

[46] H. Shin, H. Roth, M. Gao, L. Lu, Z. Xu, J. Yao, D. Mollura, and R. Summers. Deep convolutional neural networks for computer-aided detection: Cnn architectures, datasets, and transfer learning. In *ArXiv*, 2015.

[47] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[48] S. Singh, A. Gupta, and A. A. Efros. Unsupervised discovery of mid-level discriminative patches. In *European Conference on Computer Vision*, 2012.

[49] C. Sun, C. Gan, and R. Nevatia. Automatic concept discovery from parallel text and visual corpora. *ArXiv:1509.07225*, 2015.

[50] C. Sun, S. Shetty, R. Sukthankar, and R. Nevatia. Temporal localization of fine-grained actions in videos by domain transfer from web images. *ACM Multimedia*, pages 371–380, 2015.

[51] I. Sutskever, O. Vinyals, and Q. Le. Sequence to sequence learning with neural networks. *NIPS*, pages 3104–3112, 2014.

[52] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *arXiv preprint arXiv:1409.4842*, 2014.

[53] T. Tuytelaars, C. H. Lampert, M. B. Blaschko, and W. Buntine. Unsupervised object discovery: A comparison. *International Journal of Computer Vision*, 2009.

[54] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko. Simultaneous deep transfer across domains and tasks. *Proc. of ICCV*, 2015.

[55] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. *Proc. of IEEE CVPR*, pages 3156–3164, 2015.

[56] X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. *Proc. of ICCV*, 2015.

[57] M. Wigness, B. Draper, and J. Beveridge. Efficient label collection for unlabeled image datasets. *Proc. of IEEE CVPR*, 2015.

[58] J. Wu, Y. Yu, C. Huang, and K. Yu. Deep multiple instance learning for image classification and auto-annotation. *Proc. of CVPR*, pages 3460–3469, 2015.

[59] R. Wu, B. Wang, and Y. Yu. Harvesting discriminative meta objects with deep cnn features for scene classification. In *Proc. of ICCV*, 2015.

[60] J. Xiao, J. Hays, K. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, pages 3485–3492, 2010.

[61] Z. Yan, H. Zhang, R. Piramuthu, V. Jagadeesh, D. DeCoste, W. Di, and Y. Yu. Hd-cnn: Hierarchical deep convolutional neural network for large scale visual recognition. *Proc. of ICCV*, 2015.

[62] F. Yu, Y. Zhang, S. Song, A. Seff, and J. Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv:1506.03365*, 2015.

[63] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning deep features for scene recognition using places database. In *NIPS*, pages 487–495, 2014.